EXHIBIT 1

Inside Facebook's massive cyber-security system

17:20 26 October 2011 by **Jim Giles**
Magazine issue 2836. **Subscribe and save**

FACEBOOK has released details of the extraordinary security infrastructure it uses to fight off spam and other cyber-scams.

Known as the Facebook Immune System (FIS), the massive defence network appears to be successful: numbers released by the company this week show that less than 1 per cent of users experience spam. Yet it's not perfect. Researchers have built a novel attack that evaded the cyber-defences and extracted private material from real users' Facebook accounts.

A lot to look after *(Image: Eric Risberg/AP/PA)*

It took just three years for FIS to evolve from basic beginnings into an all-seeing set of algorithms that monitors every photo posted to the network, every status update– indeed, every click made by every one of the 800 million users. There are more than 25 billion of these "read and write actions" every day. At peak activity the system checks 650,000 actions a second.

"It's a big challenge," says Jim Larus, a Microsoft researcher in Redmond, Washington, who studies large networks. The only network bigger, Larus suspects, is the web itself. That makes Facebook's defence system one of the largest in existence.

It protects against scams by harnessing artificially intelligent software to detect suspicious patterns of behaviour. The system is overseen by a team of 30 people, but it can learn in real time and is able to take action without checking with a human supervisor.

One notable attack took place in April, says Tao Stein, a Facebook engineer who works on the system. It began when several users were duped into copying computer code into their browser's address bar. The code commandeered the person's Facebook account, and started sending chat messages to their friends saying things like "I just got a free iPad", along with a link where the friends could go to get their own. Friends who clicked on the link went to a site that encouraged them to paste the same code into their browsers, further spreading the plague. "Attacks like these can generate millions of messages per minute," says Stein.

Users are less likely to fall for a similar tactic when using email, because the message would probably be sent by a stranger.

But inside Facebook's network it's much more persuasive. "It's easier to exploit trust relationships in online social networks," says Justin Ma, a computer scientist at the University of California, Berkeley, who develops methods to combat email spam.

To tackle the attack, FIS generated a signature that it used to differentiate between spam and legitimate messages. This was based on the links in the spam messages, keywords like "free" and "iPad", and the IP addresses of the computers sending the messages.

But spammers can use multiple machines to switch IP addresses, and link redirection services like bit.ly can change links on the fly. So FIS checked to see which messages were being flagged as spam by users and blocked messages with similar keywords in the text. Together with other features of the message, which Facebook declined to discuss for fear of aiding spammers, the system was able to begin developing a signature to identify the spam within seconds of the attack emerging.

Facebook said this week that, thanks to FIS, less than 4 per cent of the network's messages are spam and that fewer than 1 in 200 users experience spam on any given day. "It's pretty good," says Ma, who has a Facebook account. "I'm pretty happy with the level of security."

Yet like any defence based on patterns of known behaviour, FIS is vulnerable to strategies it has not seen before. Yazan Boshmaf and colleagues at the University of British Columbia in Vancouver, Canada, have exploited this and eluded the system by creating "socialbots"– software that can pose as a human and control a Facebook account.

The bots began by sending friend requests to random users, around 1 in 5 of whom accepted. They then sent requests to the friends of the people they had connected with, and the acceptance rate jumped to almost 60 per cent. After seven weeks the team's 102 bots had made a combined 3000 friends.

Facebook's privacy settings allow users to shield personal information from public view. But because the socialbots posed as friends, they were able to extract some 46,500 email addresses and 14,500 physical addresses from users' profiles– information that could be used to launch phishing attacks or aid in identity theft.

"An attacker could do many things with this data," says Boshmaf, who will present the team's work at the Annual Computer Security Applications Conference in Orlando, Florida, next month.

A socialbot attack is yet to happen, but it's only a matter of time. Socialbots behave differently to humans that enter Facebook for the first time, in part because they have no real-world friends to connect with, and their random requests lead to an unusually high number of rejections. FIS would be able to use this pattern to recognise and block an attack of socialbots, says Stein. That would put Facebook back on top– if only until hackers release their next innovation.

Back to article

PRINT SEND SHARE